

# “Superluminal” FITS File Processing on Multiprocessors: Zero Time Endian Conversion Technique

Satoshi Eguchi<sup>1</sup>

*Astronomy Data Center, National Astronomical Observatory of Japan, 2-21-1, Osawa, Mitaka, Tokyo 181-8588, Japan*

satoshi.eguchi@nao.ac.jp

## ABSTRACT

The FITS is the standard file format in astronomy, and it has been extended to agree with astronomical needs of the day. However, astronomical datasets have been inflating year by year. In case of ALMA telescope, a  $\sim$  TB scale 4-dimensional data cube may be produced for one target. Considering that typical Internet bandwidth is a few 10 MB/s at most, the original data cubes in FITS format are hosted on a VO server, and the region which a user is interested in should be cut out and transferred to the user (Eguchi et al. 2012). The system will equip a very high-speed disk array to process a TB scale data cube in a few 10 seconds, and disk I/O speed, endian conversion and data processing one will be comparable. Hence to reduce the endian conversion time is one of issues to realize our system. In this paper, I introduce a technique named “just-in-time endian conversion”, which delays the endian conversion for each pixel just before it is really needed, to sweep out the endian conversion time; by applying this method, the FITS processing speed increases 20% for single threading, and 40% for multi-threading compared to CFITSIO. The speed-up by the method tightly relates to modern CPU architecture to improve the efficiency of instruction pipelines due to break of “causality”, a programmed instruction code sequence.

*Subject headings:* Astronomical databases: miscellaneous — Virtual observatory tools — Methods: data analysis

## 1. Introduction

The Flexible Image Transport System (FITS) is the standard data format for astronomical observed data even though they are products through calibration pipelines or otherwise. One FITS file can store multiple CCD images and photon event lists as tables, and this feature makes FITS format prevail from the radio band to the X-ray band. Especially, most archival datasets and source catalogs are provided as FITS files in these days.

The original purpose of the FITS format was to transport digital astronomical images from a computer to another with a magnetic tape

(Wells et al. 1981). There were no unified standard for computers at that time, and bit size assigned to a character and an integer was quite different from one model to another, even from the same makers. Thus the authors newly had to create a machine independent and future expandable image format for data exchange, FITS. Since then the FITS format has been repeatedly extended to agree with astronomical needs of the day (e.g., Greisen & Harten 1981; Grosbol et al. 1988).

However, we will look at the issue of astronomical data inflation, not of the format, in the years ahead; Atacama Large Millimeter/submillimeter Array (ALMA), which is the largest radio telescope built on the Chajnantor plateau in northern Chile, started observations last year. ALMA is es-

---

<sup>1</sup>Postdoctoral Fellow of Japanese Virtual Observatory Project.

timated to generate  $\sim 200$  TB observational raw data every year, and the volume of a processed 4-dimensional data cube<sup>1</sup> for one target may exceed  $\gtrsim 2$  TB (Lucas et al. 2004). Furthermore, Large Synoptic Survey Telescope (LSST), a project in 2020s, will generate 30 TB data every night<sup>2</sup>. We need a system which assists astronomers to find something interested in such big data.

Looking at such future, National Astronomical Observatory of Japan has been developing a large data providing system for ALMA utilizing the technology of Virtual Observatory (VO) to share our outputs with global astronomical communities; all processed datasets (FITS files) are hosted on a VO server, and an user can select cut-out region to download by a web-based graphical user interface (Eguchi et al. 2012, Paper I hereafter).

A prototype service is already public<sup>3</sup>, and I am working on its optimization now. The system has to process a TB scale data cube in a few 10 seconds for users' convenience, thus it is planned to equip a very high-speed disk array<sup>4</sup> and disk I/O speed and data processing one will be comparable. All the components of the system consist of Intel platform, which adopts little endian, while the FITS format does big endian. For the interactive TB size FITS file processing system, the endian conversion time is not negligible.

In this paper, I introduce a technique to make the endian conversion time apparently disappear, and to make the system much faster by multiprocessing. I describe the hardware and software configuration for evaluation in Section 2, and compare endian conversion algorithms and their performance in Section 3. In Section 4, I examine the best timing for endian conversion, and discuss the performance increase by the conversion timing in Section 5. Through the paper, I repeated measurements 100 times for each item, and adopted its sample standard deviation (a square root of unbiased variance) as  $1\text{-}\sigma$  statistical error, ignoring any systematic ones.

<sup>1</sup> $= (2\text{D Image}) \otimes (\text{Spectrum}) \otimes (\text{Polarization})$

<sup>2</sup><http://www.lsst.org/lsst/science/development>

<sup>3</sup><http://jvo.nao.ac.jp/portal/alma/>

<sup>4</sup>A system which consists of 16 striping solid state disks (SSDs) in the consumer products market effectively reaches  $\simeq 4$  GB/s read/write performance.

## 2. Configuration and Test Data

Table 1 shows the hardware and software configuration used for verification of the method. I used two types of CPUs, Intel Core i7-2600 (for Machine A) and AMD FX-8350 (for Machine B), to prevent bias due to microarchitecture. Through the paper, Intel Turbo Boost Technology (the former) and AMD Turbo CORE Technology (the latter) are disabled by BIOS for simplicity. In addition, Intel Hyper-Threading Technology (the former) is also disabled for the same reason. Thus Machine A and B are available 4 and 8 physical processors, respectively. The memory bandwidths and storage speeds were obtained the following commands: `dd if=/dev/zero of=/dev/null bs=1G count=100`, and `hdparm -t (device)`, respectively.

The same software is installed in both computers: Ubuntu 12.04.1 LTS (amd64), a Debian based 64-bit Linux, for operating system, GNU Compiler Collection (GCC) Version 4.6 for C/C++ compiler (`gcc/g++`), and CFITSIO Version 3.310 for C language FITS library (Pence 2010). I applied the `-O2 -pipe -Wall` compile options to CFITSIO and programs used in the paper. The Streaming SIMD<sup>5</sup> Extensions 2 (SSE2) codes in CFITSIO was enabled since I built the library on a 64-bit Linux<sup>6</sup>, but the SSSE3 option was disabled since

<sup>5</sup>Single Instruction/Multiple Data

<sup>6</sup>There is no way to make the `_SSE2_` macro undefined with 64-bit GCC, which switches the codes for SSE2 or otherwise in CFITSIO.




Fig. 1.— Test FITS image to evaluate parallelization efficiencies: a false color mosaic image of Carina Nebula obtained with Hubble Space Telescope, consisting of  $29,566 \times 14,321$  pixels (3.4 GB).

TABLE 1  
THE HARDWARE AND SOFTWARE INFORMATION USED FOR THE EVALUATIONS

	Machine A	Machine B
CPU	Intel Core i7-2600 (3.4 GHz)	AMD FX-8350 (4.0 GHz)
RAM	8 GB (8.80 ± 0.03 GB/s)	16 GB (22.04 ± 0.07 GB/s)
Storage	SSD (Read: 506.5 ± 0.7 MB/s)	HDD (Read: 143 ± 2 MB/s)
Operating System	Ubuntu 12.04.1 (amd64)	
C/C++ Compiler	GNU Compiler Collection Version 4.6	
FITS Library	CFITSIO Version 3.310	

NOTE.—Intel Turbo Boost and Hyper-Threading Technologies (for Machine A), AMD Turbo CORE Technology (for Machine B) are disabled through the paper. Hence 4 and 8 physical processors are available for Machine A and B, respectively.

the SSSE3 instruction set is treated as an extension in the amd64 environment.

I use a false color mosaic image of Carina Nebula obtained with Hubble Space Telescope<sup>7</sup> for test data. The image is public in Tagged Image File Format (TIFF), thus I converted it into a gray scale double precision FITS file with `convert` command provided by ImageMagick<sup>8</sup>. The size is 29,566 pixels in width and 14,321 pixels in height. The file volume is 3.4 GB (Figure 1). Through the paper, I put this FITS file on a tmpfs (Rohland 2001) mounted on `/run/shm`, to ensure that the file is always on memory for fast access. See Appendix A for the difference between tmpfs and ramdisk.

### 3. Endian Conversion Algorithms

#### 3.1. Formalism

Let  $(b_1, b_2, \dots, b_8)$  be a byte sequence of an internal expression of a 64-bit size value  $a$ . The 64-bit endian conversion of  $a$  can be expressed with a permutation  $\sigma$  as

$$a' = (b_{\sigma(1)}, b_{\sigma(2)}, \dots, b_{\sigma(8)}), \quad (1)$$

where

$$\sigma = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 8 & 7 & 6 & 5 & 4 & 3 & 2 & 1 \end{pmatrix} \quad (2)$$

in Cauchy’s two-line notation, and  $\sigma^2 = 1$  (Figure 2).

<sup>7</sup><http://hubblesite.org/newscenter/archive/releases/2007/16/image/a/>

<sup>8</sup><http://www.imagemagick.org/script/index.php>

#### 3.2. Implementation

##### 3.2.1. Byte Shuffle: Straightforward Implementation

A straightforward implementation of Eq. (1) and Eq. (2) can be written as follows:

```
uint64_t byte_shuffle(uint64_t a)
{
    unsigned char *p = (unsigned char *)&a;
    unsigned char tmp;

    tmp = p[7]; p[7] = p[0]; p[0] = tmp;
    tmp = p[6]; p[6] = p[1]; p[1] = tmp;
    tmp = p[5]; p[5] = p[2]; p[2] = tmp;
    tmp = p[4]; p[4] = p[3]; p[3] = tmp;

    return a;
}
```

I now call this method “byte shuffle”. One will find a short discussion about another implementation of byte shuffle algorithm in Appendix B.

##### 3.2.2. Bit Shift

Another implementation to perform endian conversion is to use both bit shift and logical operations:

```
uint64_t bit_shift(uint64_t a)
{
    return ((a & 0x00000000000000FFULL)
           << 56)
           | ((a & 0x000000000000FF00ULL)
           << 40)
           | ((a & 0x0000000000FF0000ULL)
```

```

        << 24)
    | ((a & 0x00000000FF000000ULL)
        << 8)
    | ((a & 0x000000FF00000000ULL)
        >> 8)
    | ((a & 0x0000FF0000000000ULL)
        >> 24)
    | ((a & 0x00FF000000000000ULL)
        >> 40)
    | ((a & 0xFF00000000000000ULL)
        >> 56);
}

```

Hereafter, I call this method “bit shift”.

### 3.2.3. BSWAP

Intel i486 and later processors have the BSWAP instruction, which converts the endian on a given 32-bit register. The instruction is extended in order to accept a 64-bit register in amd64 (Intel 2012). Furthermore, GCC Version 4.3 and later have a helper function to call the instruction, and its prototype is `uint64_t __builtin_bswap64(uint64_t x)`; Now I call endian conversions utilizing this function “BSWAP”.

### 3.2.4. SSE2

SSE2 is a set of vector instructions for Intel platform, became a part of default instruction set for amd64 environment. The endian conversion codes utilizing SSE2 can process two 64-bit values at once, and be written as follows:

```
#include <emmintrin.h>
```




Fig. 2.— The schematic diagram of a permutation operator  $\sigma$  for the endian conversion of a 64-bit value.

```

void sse2(uint64_t a[2])
{
    __m128i r0 = _mm_load_si128((__m128i *)a);
    // r0 <- a
    __m128i r1 = _mm_srli_epi16(r0, 8);
    // 8-bit shifts towards right
    // for four 2-byte integers
    __m128i r2 = _mm_slli_epi16(r0, 8);
    // 8-bit shifts towards left
    // for four 2-byte integers
    r0 = _mm_or_si128(r1, r2);
    // 128-bit or operation
    // on r1 and r2
    r0 = _mm_shufflelo_epi16(r0,
        _MM_SHUFFLE(0, 1, 2, 3));
    // byte shuffle for the
    // lower half of r0 register
    r0 = _mm_shufflehi_epi16(r0,
        _MM_SHUFFLE(0, 1, 2, 3));
    // byte shuffle for the
    // higher half of r0 register
    _mm_store_si128((__m128i *)a, r0);
    // a <- r0
}

```

There are almost the same codes in CFITSIO and SLLIB/SFITSIO<sup>9</sup>. I call these codes simply “SSE2”, hereafter.

### 3.3. SSSE3

Another vector instruction set called “SSSE3” is available for Intel Core series and later CPUs. Utilizing this instruction set, one can perform endian conversion of two 64-bit values at one instruction. An example is follows:

```

#include <tmmintrin.h>
void ssse3(uint64_t a[2])
{
    static const __m128i mask
        = _mm_set_epi8(
            8, 9, 10, 11, 12, 13, 14, 15,
            0, 1, 2, 3, 4, 5, 6, 7
        );
    __m128i r = _mm_load_si128((__m128i *)a);
    __m128i r = _mm_shuffle_epi8(r, mask);
    _mm_store_si128((__m128i *)a, r);
}

```

<sup>9</sup><http://www.ir.isas.jaxa.jp/~cyamauch/sli/index.html>

TABLE 2  
ENDIAN CONVERSION TIME

Machine	Bit Shift (msec)	BSWAP (msec)	SSE2 (msec)	SSSE3 (msec)	Byte Shuffle (msec)
Machine A	410 $\pm$ 2	410 $\pm$ 2	405 $\pm$ 2	372.4 $\pm$ 0.1	3190.3 $\pm$ 0.6
Machine B	601.3 $\pm$ 0.6	605.8 $\pm$ 0.7	582.3 $\pm$ 0.7	598 $\pm$ 2	8056.3 $\pm$ 0.3

NOTE.—The endian conversion time of 423,414,686 ( $=29,566 \times 14,321$ ) double-type elements with various algorithms.

There are almost same codes in CFITSIO too. I call these codes simply “SSSE3”, hereafter.

### 3.4. Benchmark

To see which one is fastest and how they behave towards parallelization, I performed simple benchmark. In the benchmark, I reserved a double-type array whose number of elements were set to  $29,566 \times 14,321 = 423,414,686$ , just the number of pixels in Figure 1, and filled the array for uniform real random numbers of 32-bit resolution on  $[-1000, 1000]$  generated with Mersenne Twister (Matsumoto & Nishimura 1998).

#### 3.4.1. Single Thread

The results are summarized in Table 2. For Machine A, three algorithms except for SSSE3 and byte shuffle process the test data in about 410 milliseconds, while SSSE3 does about 370 millisecond. On the other hand, for Machine B, four algorithms except for byte shuffle process the test data in about 600 milliseconds, and the SSE2 algorithm is fastest in the all ones. The byte shuffle algorithm is slowest by one order compared to the others.

#### 3.4.2. Multi-Thread

I also examined the CPU-scalability of these algorithms. I adopted `pthread` for parallelization, and simply divided the array containing the test data into equal-size segments so that the total number of the segments were equal to the number of threads. Then I assigned each thread with each segment.

Figure 3 represents the results. I also list the observed values for detailed comparison of the algorithms in Table 3 (for Machine A) and Table 4 (for Machine B). Except for byte shuffle algorithm,

I observed  $\simeq 10\%$  performance gain for Machine A, and  $\simeq 40\%$  up to four threads for Machine B with the four algorithms.

It seems strange that the memory bandwidth of Machine B is sufficient for the test data size but the four algorithms show performance cutoff at four threads. I performed detailed hardware benchmark utilizing Lmbench<sup>10</sup>, and found that context switching time and the latency of L2 cache memory normalized in CPU cycles of Machine B are 2.4 times and 4.6 times, respectively, larger than those of Machine A. Hence I conclude that there are some hardware bottlenecks in Machine B, which cause the plateau in Figure 3.

The behaviors of the four algorithms with respect to the number of threads are very similar, and I adopt bit shift algorithm in the next section because of its compiler portability and identicalness to BSWAP (see Appendix C).

## 4. Endian Conversion Timing

A modern CPU has multiple arithmetic logic units (ALUs) and instruction pipelines to boost the operating rates of ALUs. As seen in the previous section, the hardware limitation lies just below the endian conversion time of single thread (Figure 3, Machine A), preventing the CPU scalability. This may lead to many holes (or “no operation” instructions) in the pipelines and reduce the performance. If this is the case, shuffling instructions in source codes can produce improvement.

To verify this assumption, I disabled the endian conversion functionality in CFITSIO; I changed the `BYTESWAPPED` macros for i386 and amd64 architectures from `TRUE` into `FALSE` in `fitsio2.h`, and commented out the codes which CFITSIO

<sup>10</sup><http://www.bitmover.com/lmbench/>

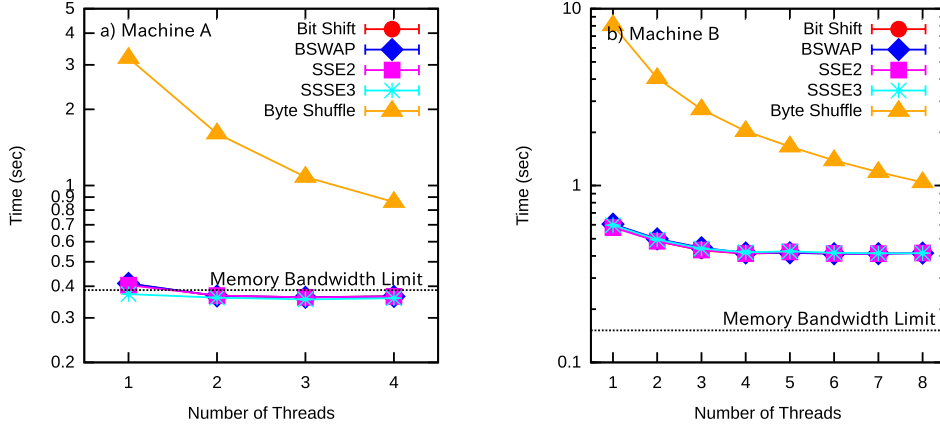


Fig. 3.— CPU-scalability comparison between the endian conversion algorithms. All the algorithms except for byte shuffle seem to behave in the same way and scale up to only the half number of CPU cores due to the hardware I/O limits.

TABLE 3  
THE CPU-SCALABILITY OF THE ENDIAN CONVERSION ALGORITHMS ON MACHINE A

Number of Threads	Bit Shift (msec)	BSWAP (msec)	SSE2 (msec)	SSSE3 (msec)	Byte Shuffle (msec)
1	410 $\pm$ 2	411 $\pm$ 3	404 $\pm$ 2	372.4 $\pm$ 0.2	3191.8 $\pm$ 0.7
2	366 $\pm$ 3	366 $\pm$ 3	367 $\pm$ 2	360.4 $\pm$ 0.5	1605 $\pm$ 2
3	362 $\pm$ 3	362 $\pm$ 3	362 $\pm$ 3	354.9 $\pm$ 0.3	1082 $\pm$ 6
4	365 $\pm$ 4	365 $\pm$ 4	365 $\pm$ 4	358.8 $\pm$ 0.5	860 $\pm$ 20

NOTE.—The CPU-scalability of 423,414,686 ( $=29,566 \times 14,321$ ) double-type element endian conversion with various algorithms on Machine A. Different from Table 2, 16-byte memory alignment is adopted.

TABLE 4  
THE CPU-SCALABILITY OF THE ENDIAN CONVERSION ALGORITHMS ON MACHINE B

Number of Threads	Bit Shift (msec)	BSWAP (msec)	SSE2 (msec)	SSSE3 (msec)	Byte Shuffle (msec)
1	597.8 $\pm$ 0.7	605.1 $\pm$ 0.7	579.3 $\pm$ 0.9	597 $\pm$ 2	8056.2 $\pm$ 0.2
2	484 $\pm$ 1	499 $\pm$ 1	484 $\pm$ 1	495 $\pm$ 1	4046 $\pm$ 3
3	429.9 $\pm$ 0.8	445 $\pm$ 3	432 $\pm$ 1	440 $\pm$ 1	2699 $\pm$ 4
4	413 $\pm$ 3	412 $\pm$ 2	413 $\pm$ 3	419 $\pm$ 7	2031 $\pm$ 8
5	418 $\pm$ 3	417 $\pm$ 1	422 $\pm$ 2	424 $\pm$ 2	1661 $\pm$ 5
6	412 $\pm$ 1	411.5 $\pm$ 0.8	413 $\pm$ 1	416 $\pm$ 1	1388 $\pm$ 4
7	412.3 $\pm$ 0.9	412.2 $\pm$ 0.5	411.9 $\pm$ 0.6	414.1 $\pm$ 0.5	1192 $\pm$ 4
8	416.1 $\pm$ 0.6	415.1 $\pm$ 0.7	415.5 $\pm$ 0.4	414.5 $\pm$ 0.4	1043 $\pm$ 2

NOTE.—All conditions are same as Table 3.

perform runtime check to verify whether the machine endian definition by the above macro is consistent with the execution environment in `cfileio.c`, and I rebuilt the library. The patches for those files are shown in Appendix E.

I compare the following two methods;

1. It loads the full test image (Figure 1) from `tmpfs` (see §2) onto an array, then it converts the endian by the parallelized bit shift algorithm (described in §3.4.2), and it sums up all the elements.
2. It loads the full test image onto an array, and it sums up all the elements with converting the endian one after another by the bit shift algorithm.

From here, I refer to the former as “on ahead endian conversion method”, and to the latter as “just-in-time endian conversion method”.

On ahead endian conversion method can be written as follows:

```
{
    double *v;    // an array to store
                  // a FITS image
    size_t len;  // the length of
                  // the array v

    // load a byte sequence from a FITS
    // file into v here...

    // endian conversion
    for (size_t i = 0; i < len; ++i) {
        uint64_t *p = (uint64_t *)&v[i];
        uint64_t a = bit_shift(*p);
        double *q = (double *)&a;
        v[i] = *q;
    }

    // process v here...
}
```

and just-in-time endian conversion method can be written as follows:

```
{
    double *v;    // an array to store
                  // a FITS image
    size_t len;  // the length of
                  // the array v
```

```
    // load a byte sequence from a FITS
    // file into v here...

    // image processing...
    {
        // something...

        // one needs to refer the value
        // of v[i] here

        // endian conversion
        uint64_t *p = (uint64_t *)&v[i];
        uint64_t a = bit_shift(*p);
        double *q = (double *)&a;
        double x = *q;
        // use x instead of v[i] below

        // something...
    }
}
```

where `bit_shift()` is the endian conversion function defined in §3.2.2.

In this section, I adopt summing up all the elements in the test image as an example of image processing.

#### 4.1. Single Thread

I implemented both methods in single thread and performed benchmark. The codes of on ahead conversion method are following:

```
{
    // endian conversion
    for (size_t i = 0; i < len; ++i) {
        uint64_t *p = (uint64_t *)&v[i];
        uint64_t a = bit_shift(*p);
        double *q = (double *)&a;
        v[i] = *q;
    }

    // summation
    double sum = 0.0;
    for (size_t i = 0; i < len; ++i) {
        sum += v[i];
    }
}
```

and those of just-in-time endian conversion method are following:

```

{
    double sum = 0.0;
    for (size_t i = 0; i < len; ++i) {
        // endian conversion
        uint64_t *p = (uint64_t *)&v[i];
        uint64_t a = bit_shift(*p);
        double *q = (double *)&a;

        // summation
        sum += *q;
    }
}

```

Note that the former codes are identical to those with original CFITSIO.

The results are summarized in Table 5. I obtained slightly faster ( $\simeq 5\%$ ) total processing time of  $2.22 \pm 0.04$  sec and  $3.62 \pm 0.04$  sec for Machine A and B, respectively, with on ahead endian conversion method, while that with original CFITSIO is  $2.38 \pm 0.04$  and  $3.79 \pm 0.03$  for Machine A and B, respectively.

On the other hand, I obtained significantly faster time of  $1.85 \pm 0.04$  and  $3.05 \pm 0.03$  for Machine A and B, respectively, which corresponds to  $\simeq 25\%$  performance gain, with just-in-time endian conversion method.

#### 4.2. Multi-Thread

I made both methods multithreaded by utilizing OpenMP<sup>11</sup> APIs for its simple implementation. The codes of just-in-time endian conversion method, for example, are below:

```

{
    double sum = 0.0;
    #pragma omp parallel for reduction (+:sum)\\
    schedule (auto)
    for (size_t i = 0; i < len; ++i) {
        // endian conversion
        uint64_t *p = (uint64_t *)&v[i];
        uint64_t a = bit_shift(*p);
        double *q = (double *)&a;

        // summation
        sum += *q;
    }
}

```

<sup>11</sup><http://openmp.org/wp/>

On the other hand, I could not find the best parameters in OpenMP APIs for the endian conversion routine in on ahead conversion method, hence I applied OpenMP only to the summation routine, and adopted the `pthread`-based parallelization described in §3.4.2 for the endian conversion routine in on ahead conversion method; the number of the threads for OpenMP was set to that for the endian conversion.

The results obtained with these programs are summarized in Table 6 (for Machine A), Table 7 (for Machine B), Figure 4 (for on ahead endian conversion method), and Figure 5 (for just-in-time endian conversion method). Note that the endian conversion time of on ahead endian conversion method is included in the FITS reading time. The total time to perform the same things with the original CFITSIO in single thread is superimposed on these figures as a dotted line:  $2.38 \pm 0.04$  seconds for Machine A, and  $3.79 \pm 0.03$  seconds for Machine B.

For the on ahead endian conversion method, the total time slightly scales the number of threads and gets faster than original CFITSIO, while the file reading time (including endian conversion time) seems to be little scalable. The scalability of the total time mostly owes that of the summation routine, and the parallelization of the endian conversion has little impact due to the hardware limit seen in §3.4.2.

On the other hand, for the just-in-time endian conversion method, the total time is interestingly smaller than that of original CFITSIO even for single thread. The summation routine seems to be scalable almost in the full range, while the total time scales up to four threads.

## 5. Discussion

### 5.1. Performance Analysis of the Simple Summation Codes

There is a well-known equation to estimate the increase by parallelization, Amdahl's law (Amdahl 1967):

$$T_{\text{parallel}} = \left\{ (1 - P) + \frac{P}{N} + \alpha \right\} T_{\text{single}}, \quad (3)$$

where  $T_{\text{single}}$  and  $T_{\text{parallel}}$  represent processing time in single thread and multi-thread cases, re-



TABLE 5  
THE DATA PROCESSING TIMES WITH TWO DIFFERENT METHOD IN SINGLE THREAD

Method	Machine	FITS Read Time (sec)	Sum Up Time (sec)	Total Time (sec)
On Ahead Endian Conversion	Machine A	$1.006 \pm 0.003$	$0.4198 \pm 0.0002$	$2.22 \pm 0.04$
	Machine B	$1.859 \pm 0.007$	$0.557 \pm 0.001$	$3.62 \pm 0.04$
Lasy Endian Conversion	Machine A	$1.010 \pm 0.003$	$0.44003 \pm 0.00008$	$1.85 \pm 0.04$
	Machine B	$1.874 \pm 0.007$	$0.621 \pm 0.002$	$3.06 \pm 0.03$

NOTE.—The total time with original CFITSIO is  $2.38 \pm 0.04$  and  $3.79 \pm 0.03$  for Machine A and B, respectively.

TABLE 6  
THE CPU-SCALABILITY OF ON AHEAD AND JUST-IN-TIME ENDIAN CONVERSION METHODS ON MACHINE  
A

Method	Number of Threads	FITS Read Time (sec)	Sum Up Time (sec)	Total Time (sec)
On Ahead Endian Conversion Method	1	$1.445 \pm 0.003$	$0.4205 \pm 0.0001$	$2.38 \pm 0.04$
	2	$1.404 \pm 0.003$	$0.220 \pm 0.004$	$2.1 \pm 0.1$
	3	$1.401 \pm 0.003$	$0.181 \pm 0.002$	$2.11 \pm 0.09$
	4	$1.403 \pm 0.003$	$0.177 \pm 0.002$	$2.09 \pm 0.03$
Just-in-Time Endian Conversion Method	1	$1.048 \pm 0.003$	$0.4416 \pm 0.0009$	$2.0 \pm 0.1$
	2	$1.051 \pm 0.003$	$0.2262 \pm 0.0005$	$1.8 \pm 0.1$
	3	$1.051 \pm 0.003$	$0.183 \pm 0.001$	$1.77 \pm 0.02$
	4	$1.052 \pm 0.003$	$0.177 \pm 0.002$	$1.76 \pm 0.02$

NOTE.—The endian conversion time is included in the FITS reading time for on ahead endian conversion method.

TABLE 7  
THE CPU-SCALABILITY OF ON AHEAD AND JUST-IN-TIME ENDIAN CONVERSION METHODS ON MACHINE  
B

Method	Number of Threads	FITS Read Time (sec)	Sum Up Time (sec)	Total Time (sec)
On Ahead Endian Conversion Method	1	$2.526 \pm 0.005$	$0.5522 \pm 0.0006$	$3.76 \pm 0.03$
	2	$2.403 \pm 0.005$	$0.299 \pm 0.003$	$3.39 \pm 0.01$
	3	$2.347 \pm 0.005$	$0.220 \pm 0.002$	$3.253 \pm 0.010$
	4	$2.326 \pm 0.005$	$0.193 \pm 0.003$	$3.203 \pm 0.008$
	5	$2.334 \pm 0.006$	$0.205 \pm 0.003$	$3.224 \pm 0.010$
	6	$2.326 \pm 0.005$	$0.189 \pm 0.003$	$3.20 \pm 0.01$
	7	$2.325 \pm 0.006$	$0.178 \pm 0.002$	$3.19 \pm 0.01$
	8	$2.334 \pm 0.005$	$0.174 \pm 0.003$	$3.19 \pm 0.01$
Just-in-Time Endian Conversion Method	1	$1.914 \pm 0.003$	$0.582 \pm 0.002$	$3.19 \pm 0.05$
	2	$1.917 \pm 0.003$	$0.328 \pm 0.002$	$2.934 \pm 0.009$
	3	$1.916 \pm 0.004$	$0.242 \pm 0.002$	$2.853 \pm 0.008$
	4	$1.916 \pm 0.004$	$0.205 \pm 0.002$	$2.812 \pm 0.009$
	5	$1.918 \pm 0.003$	$0.205 \pm 0.001$	$2.816 \pm 0.009$
	6	$1.918 \pm 0.003$	$0.1948 \pm 0.0009$	$2.801 \pm 0.010$
	7	$1.917 \pm 0.004$	$0.185 \pm 0.001$	$2.791 \pm 0.010$
	8	$1.917 \pm 0.003$	$0.178 \pm 0.003$	$2.781 \pm 0.010$

NOTE.—The endian conversion time is included in the FITS reading time for on ahead endian conversion method.

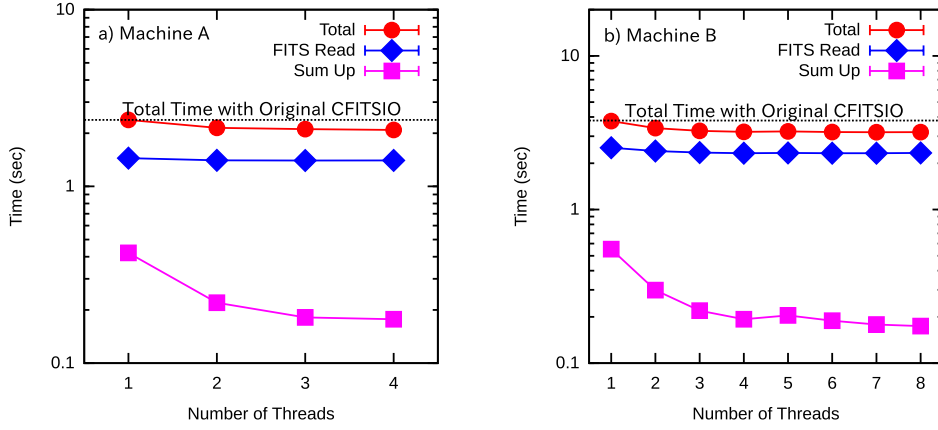


Fig. 4.— The CPU-scalability of the FITS reading and processing time by applying the bit shift parallel algorithm to CFITSIO. I also applied a simple OpenMP parallelization to the summation routine.

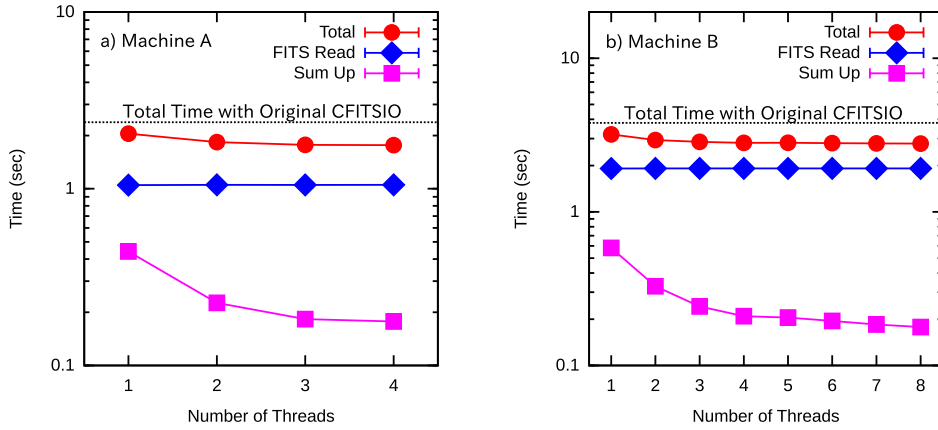


Fig. 5.— The CPU-scalability of the just-in-time endian conversion algorithm. Endian conversions are performed in the summation routine, which is applied a simple OpenMP parallelization.

spectively,  $P$  is the ratio of codes which parallelization methods are applied to<sup>12</sup>,  $N$  is the number of threads, and  $\alpha$  is the overhead caused by parallelization.

To quantify the performance increase of on ahead endian conversion method and just-in-time endian conversion method, I performed model fitting to the total time of both methods with Eq.(3). I found that  $\alpha \sim O(10^{-30})$  while the fitting, thus I fixed  $\alpha$  at 0. The results are summarized in Table 8 and Figure 6. The increasing rates of performance compared to original CFITSIO ( $T_{\text{single}} = 2.38 \pm 0.04$  for Machine A and  $T_{\text{single}} = 3.79 \pm 0.03$  for Machine B) are also listed in the table.

The figure shows that the above results are explained well by Amdahl's law, and that the on ahead endian conversion method for single thread has almost the same performance as original CFITSIO. In fact, these two agree with each other in  $\lesssim 5\%$  errors according to the table. The table also suggests that multi-threading boosts this method up about 20%. Considering the parallelization rate  $P \simeq 16\%$ , one cannot expect further speed up by multi-threading in  $N \gtrsim 4$ . This suggests that the bottlenecks of other hardware disrupt order in the instruction pipelines and leads to the decrease of operating ratio of ALUs.

On the other hand, the just-in-time endian conversion method is 20% faster than both of original CFITSIO and the single thread version of on ahead one, surprisingly. This seems as if the endian conversion process disappeared. In the parallelized case, the just-in-time conversion method is 40% faster than the others in single thread. However, the performance increase by multi-threading can be expected only in  $N \lesssim 4$  since the parallelization rate  $P \simeq 16\%$ , due to the hardware bottlenecks mentioned above.

For further investigation, I fitted the summation time of these methods with Eq.(3) to investigate the impact of the endian conversion codes in the summation routine on performance; there are endian conversion codes in the summation routine in case of just-in-time endian conversion method, but not in case of on ahead endian conversion method. The results are summarized in Table 9 and Figure 7. I found that the parallelization rate  $P \approx 85\%$  in both cases, and

that the ratio of  $T_{\text{single}}$  of just-in-time endian conversion method against that of on ahead one  $r = T_{\text{single}}(\text{Just-in-Time}) / T_{\text{single}}(\text{On Ahead})$  was equal to  $r = 1.02 \pm 0.05$  for Machine A and  $r = 1.03 \pm 0.04$  for Machine B. There is no overhead of endian conversion in the summation routine, since the shift of  $r$  from unity is not significant statistically.

Thus I conclude that endian conversion is so simple operation for a modern CPU that the bottlenecks of other hardware disrupt order in the instruction pipelines; to prevent the disruption, the endian conversion should be done just before a value is referred.

## 5.2. Application to ALMAWebQL

From here, I only investigated the performance increase of summing up all the elements in a large FITS file by just-in-time endian conversion method. In this subsection, I apply the method to ALMAWebQL, our interactive web viewer for ALMA data cubes described in Paper I, to obtain more realistic benchmark data. For realistic and fair comparison, the SSE2 boosted endian conversion codes in CFITSIO are enabled for on ahead endian conversion method, while there is no SSE2 code in just-in-time endian conversion method.

ALMA data cubes not contain information of polarization currently, and they are simple 3-dimensional FITS files (Figure 8). For image extraction, one have to integrate the cube along the spectral direction; for spectrum extraction, one convolute all spatial information.




Fig. 8.— The schematic illustration of a data cube of ALMA Science Verification Data.

<sup>12</sup>Hardware bottlenecks are included in the  $1 - P$  term.

TABLE 8  
THE FITTING RESULTS OF THE TOTAL PROCESSING TIME

Method	Machine	$T_{\text{single}}$ (sec)	$P$	$\chi^2$ (d.o.f. <sup>a</sup> )	Increase Rate of Performance	
					Single Thread	Multi-Thread
On Ahead Endian Conversion	Machine A	$2.377 \pm 0.009$	$0.164 \pm 0.006$	0.09 (2)	$1.00 \pm 0.02$	$1.20 \pm 0.02$
	Machine B	$3.68 \pm 0.05$	$0.16 \pm 0.02$	49 <sup>b</sup> (6)	$1.03 \pm 0.02$	$1.22 \pm 0.03$
Just-in-Time Endian Conversion	Machine A	$2.02 \pm 0.05$	$0.17 \pm 0.03$	0.5 (2)	$1.18 \pm 0.04$	$1.42 \pm 0.07$
	Machine B	$3.13 \pm 0.02$	$0.130 \pm 0.009$	8.5 (6)	$1.21 \pm 0.01$	$1.39 \pm 0.02$

NOTE.—The fitting results of the total processing time with respect to two different endian conversion methods with Amdahl's law and their increase rate of performance compared with original CFITSIO ( $T_{\text{single}} = 2.38 \pm 0.04$  for Machine A and  $T_{\text{single}} = 3.79 \pm 0.03$  for Machine B). The errors are 1- $\sigma$  confidence limits for a single parameter.

<sup>a</sup>Degrees of freedom

<sup>b</sup>The large  $\chi^2$  value is caused by very small errors of observed values, which agree with the model well (see Figure 6).

TABLE 9  
THE FITTING RESULTS OF THE TIME TO SUM UP ALL ELEMENTS

Method	Machine	$T_{\text{single}}$ (msec)	$P$	$\chi^2$ (d.o.f.)
On Ahead Endian Conversion	Machine A	$420 \pm 1$	$0.82 \pm 0.04$	127.9 (2)
	Machine B	$551 \pm 6$	$0.81 \pm 0.02$	491.1 (6)
Just-in-Time Endian Conversion	Machine A	$430 \pm 20$	$0.92 \pm 0.06$	913.4 (2)
	Machine B	$570 \pm 20$	$0.80 \pm 0.02$	312.7 (6)

NOTE.— $T_{\text{single}}$  (Just – in – Time) /  $T_{\text{single}}$  (On Ahead) =  $1.02 \pm 0.05$  (for Machine A),  $1.03 \pm 0.04$  (for Machine B).

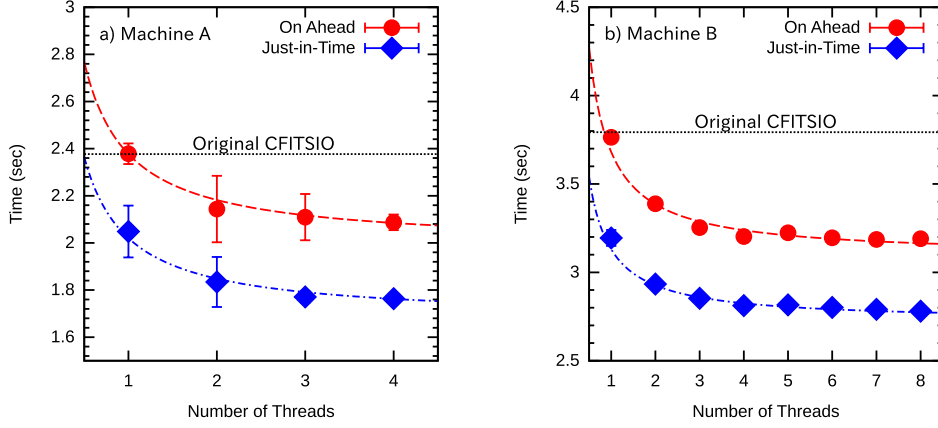


Fig. 6.— The fitting results of on ahead and just-in-time endian conversion method with respect to total processing time with Amdahl's law. The red dashed and blue dash dotted lines correspond to the law for on ahead and just-in-time endian conversion method, respectively.

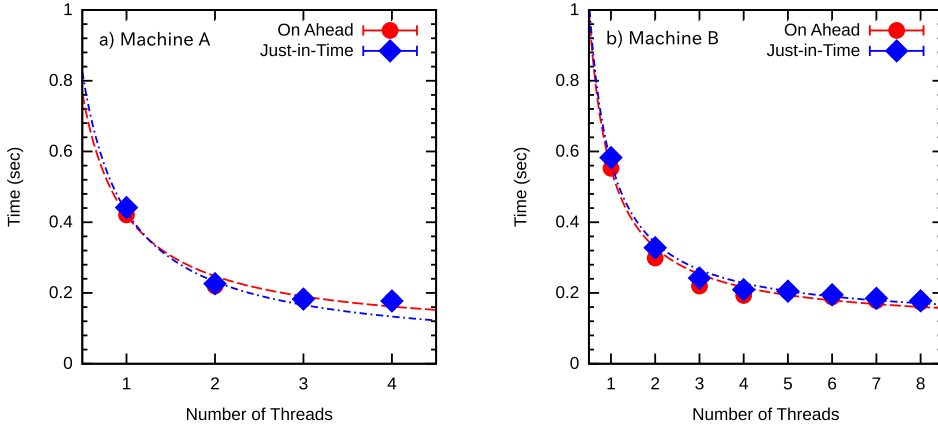


Fig. 7.— The fitting results of on ahead and just-in-time endian conversion method with respect to summation time with Amdahl's law. The red dashed and blue dash dotted lines correspond to the law for on ahead and just-in-time endian conversion method, respectively.

TABLE 10  
IMAGE EXTRACTION TIME

File Size (MB)	On Ahead Endian Conversion Method	Just-in-Time Endian Conversion Method
21.4	$0.02953 \pm 0.00006$	$0.0361 \pm 0.0004$
55.0	$0.75214 \pm 0.00003$	$0.0910 \pm 0.0001$
66.9	$0.10007 \pm 0.00006$	$0.10295 \pm 0.00008$
78.4	$0.13136 \pm 0.00008$	$0.12371 \pm 0.00008$
88.7	$0.14435 \pm 0.00007$	$0.1299 \pm 0.0001$
106.9	$0.1835 \pm 0.0001$	$0.1606 \pm 0.0002$
150.3	$0.2333 \pm 0.0001$	$0.21006 \pm 0.00008$

NOTE.—The time to extract an image from an ALMA data cube on Machine A in single thread.



TABLE 11  
SPECTRUM EXTRACTION TIME

File Size (MB)	On Ahead Endian Conversion Method	Just-in-Time Endian Conversion Method
21.4	$0.02743 \pm 0.00002$	$0.0344 \pm 0.0004$
55.0	$0.06970 \pm 0.00002$	$0.07527 \pm 0.00008$
66.9	$0.08512 \pm 0.00002$	$0.08583 \pm 0.00010$
78.4	$0.09861 \pm 0.00003$	$0.0909 \pm 0.0001$
88.7	$0.11271 \pm 0.00003$	$0.10539 \pm 0.00008$
106.9	$0.13394 \pm 0.00004$	$0.1153 \pm 0.0003$
150.3	$0.18769 \pm 0.00005$	$0.15503 \pm 0.00006$

NOTE.—The time to extract a spectrum from an ALMA data cube on Machine A in single thread.

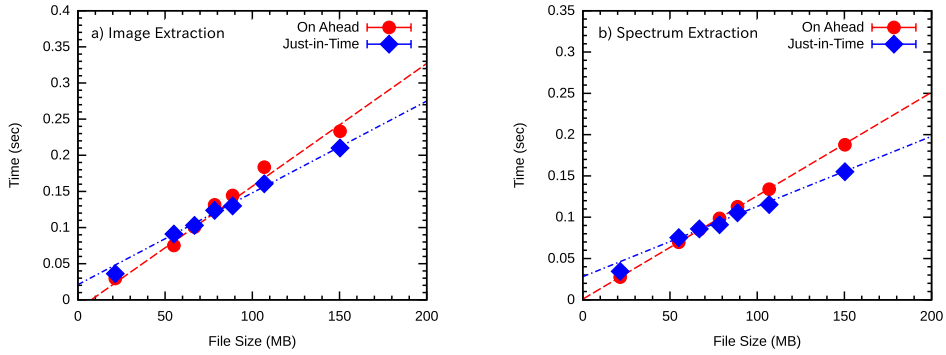


Fig. 9.— The image (left) and spectrum (right) extraction time with respect to file size of ALMA data cube. The red dashed and blue dash dotted lines correspond to the best fit ones for on ahead and just-in-time endian conversion method, respectively.

I measured the time to complete these computations in single thread with various size data on Machine A. The results for image extraction are summarized in Table 10, and those for spectrum extraction are summarized in Table 11. From these tables, I obtain

$$T(\text{On Ahead}) = (1.7 \pm 0.1) \times 10^{-3} \left( \frac{V}{\text{MB}} \right) - (0.013 \pm 0.007) \text{ sec}, \quad (4)$$

$$T(\text{Just-in-Time}) = (1.27 \pm 0.04) \times 10^{-3} \left( \frac{V}{\text{MB}} \right) + (0.021 \pm 0.003) \text{ sec} \quad (5)$$

for image extraction and

$$T(\text{On Ahead}) = (1.252 \pm 0.007) \times 10^{-3} \left( \frac{V}{\text{MB}} \right) + (0.0008 \pm 0.0004) \text{ sec}, \quad (6)$$

$$T(\text{Just-in-Time}) = (0.85 \pm 0.02) \times 10^{-3} \left( \frac{V}{\text{MB}} \right) + (0.028 \pm 0.002) \text{ sec} \quad (7)$$

for spectrum extraction, where  $T(\text{On Ahead})$  and  $T(\text{Just-in-Time})$  represent the time with on ahead and just-in-time endian conversion methods, respectively, and  $V$  is file size in the MB unit (Figure 9). Hence just-in-time endian conversion method in single thread is  $\gtrsim 20\%$  faster than on ahead conversion method boosted by SSE2 above  $V \gtrsim 200$  MB. This demonstrates that just-in-time endian conversion method can be very powerful when one performs convolution and stacking of very large images, which are very common analysis techniques in optical band, obtained with future large telescopes.

### 5.3. Data Types

In this paper, I only treated a double precision FITS file, but one could expect almost the same results for float and LONG data types, which correspond to BITPIX = -32 and 32, respectively; as demonstrated in Appendix C, `bit_shift()` function is compiled into BSWAP instruction. The amd64 architecture can handle both of 32 bit and 64 bit operation codes and their operands seamlessly. On the other hand, for byte and short data

type, there may be little advantage of just-in-time endian conversion method since BSWAP instruction cannot take any 16 bit values as its operand, and up-casting into 32 bit integer always occurs in arithmetic operations in both cases.

## 6. Summary

The FITS format was originally developed to exchange digital astronomical datasets from a computer to another, but the progress of computation power and software technology enables one to process FITS files through web browsers. In addition, data size has been inflating year by year, and it will exceed  $\sim$  TB in the year ahead. To handle such big FITS file with web applications, the endian conversion time from the FITS native to the machine one cannot be negligible, and a solution for this problem is required.

In this paper, I compared the features of four typical endian conversion algorithms under multi-thread environment, and found the bit shift one was suitable for parallelization. Then I examined the best timing for endian conversion under multi-thread environment. I found that one should postpone the endian conversion until a value is really referred in a program, because endian conversion is so simple for a modern CPU that the bottlenecks of other hardwares disrupt order in the instruction pipelines, which leads to the decrease of operating ratio of ALUs. In fact, by applying this method to loading 3.4 GB FITS file and sum up all the elements, the performance increased 20% for single thread and 40% for multi-thread compared to CFITSIO, which corresponded to  $\gtrsim 600$  milliseconds, and one can be aware of the speed-up. No overhead of endian conversion was found on the summation routine; hence one can sweep the endian conversion time out of his/her codes. Note that parallelization of this method peaked out in four threads in the experiment.

CPU vendors introduce various techniques, such as speculative execution and branch prediction, to improve the efficiency of instruction pipelines; an executed instruction code sequence is apart from a programmed one. In this context, modern CPUs partially break “causality”, a programmed instruction code sequence, and gain speed. Just-in-time endian conversion method utilizes such boosting technology. There is nothing

new in the method, but it must be a small step to handle astronomical big data generated by the next generation telescopes.

I greatly appreciate Dr. Chisato Yamauchi, who is my colleague and the author of SLIB/SFITSIO<sup>13</sup>, for rewarding discussions.

---

<sup>13</sup>SFITSIO is a light weight FITS library for C/C++, providing modern APIs.

## A. Tmpfs and Ramdisk

Both tmpfs and ramdisk are a data space allocated on memory. One has to specify the size in advance for ramdisk, while one does not set the size for tmpfs in advance necessarily since it is under control of virtual memory manager and shares swap space.

When an application requests the operating system for memory blocks and when there does not remain sufficient physical memory space, the memory manager firstly swap out the files on tmpfs. Tmpfs is ideal space to put temporal files which one requires very fast access to.

## B. Another Implementation of Byte Shuffle Algorithm

One can also implements byte shuffle algorithm as follows:

```
uint64_t byte_shuffle2(uint64_t a)
{
    unsigned char *p = (unsigned char *)&a;
    uint64_t b;
    unsigned char *q = (unsigned char *)&b;

    q[0] = p[7];
    q[1] = p[6];
    q[2] = p[5];
    q[3] = p[4];
    q[4] = p[3];
    q[5] = p[2];
    q[6] = p[1];
    q[7] = p[0];

    return b;
}
```

The number of assignments of the codes (= 8) is less than that shown in the main part of this paper (= 12), and one would expect further performance improvement.

I disassembled both two codes compiled with the -O2 option, and obtained followings:

```
0000000000000000 <byte_shuffle>:
  0: 49 89 fa      mov     %rdi,%r10
  3: 49 89 f8      mov     %rdi,%r8
  6: 89 fe         mov     %edi,%esi
  8: 48 89 f9      mov     %rdi,%rcx
  b: 89 fa         mov     %edi,%edx
  d: 48 89 f8      mov     %rdi,%rax
 10: 40 88 7c 24 ff  mov     %dil,-0x1(%rsp)
 15: 48 c1 e8 20    shr     $0x20,%rax
 19: 49 c1 ea 38    shr     $0x38,%r10
 1d: 49 c1 e8 30    shr     $0x30,%r8
 21: 66 c1 ee 08    shr     $0x8,%si
 25: 48 c1 e9 28    shr     $0x28,%rcx
 29: c1 ea 10      shr     $0x10,%edx
 2c: c1 ef 18      shr     $0x18,%edi
 2f: 44 88 54 24 f8  mov     %r10b,-0x8(%rsp)
```

```

34: 40 88 74 24 fe      mov     %sil,-0x2(%rsp)
39: 44 88 44 24 f9      mov     %r8b,-0x7(%rsp)
3e: 88 54 24 fd          mov     %dl,-0x3(%rsp)
42: 88 4c 24 fa          mov     %cl,-0x6(%rsp)
46: 40 88 7c 24 fc      mov     %dil,-0x4(%rsp)
4b: 88 44 24 fb          mov     %al,-0x5(%rsp)
4f: 48 8b 44 24 f8      mov     -0x8(%rsp),%rax
54: c3                  retq

```

, and

0000000000000000 <byte\_shuffle2>:

```

0: 49 89 fa            mov     %rdi,%r10
3: 49 89 f9            mov     %rdi,%r9
6: 49 89 f8            mov     %rdi,%r8
9: 48 89 fe            mov     %rdi,%rsi
c: 89 f9              mov     %edi,%ecx
e: 89 fa              mov     %edi,%edx
10: 89 f8              mov     %edi,%eax
12: 49 c1 ea 38         shr     $0x38,%r10
16: 49 c1 e9 30         shr     $0x30,%r9
1a: 66 c1 e8 08         shr     $0x8,%ax
1e: 49 c1 e8 28         shr     $0x28,%r8
22: 48 c1 ee 20         shr     $0x20,%rsi
26: c1 e9 18           shr     $0x18,%ecx
29: c1 ea 10           shr     $0x10,%edx
2c: 44 88 54 24 f8      mov     %r10b,-0x8(%rsp)
31: 44 88 4c 24 f9      mov     %r9b,-0x7(%rsp)
36: 44 88 44 24 fa      mov     %r8b,-0x6(%rsp)
3b: 40 88 74 24 fb      mov     %sil,-0x5(%rsp)
40: 88 4c 24 fc          mov     %cl,-0x4(%rsp)
44: 88 54 24 fd          mov     %dl,-0x3(%rsp)
48: 88 44 24 fe          mov     %al,-0x2(%rsp)
4c: 40 88 7c 24 ff      mov     %dil,-0x1(%rsp)
51: 48 8b 44 24 f8      mov     -0x8(%rsp),%rax
56: c3                  retq

```

, that is, there are less assignments in `byte_shuffle2()` (= 8) than `byte_shuffle()` (= 12), however, the former binary codes are longer than the latter ones. Hence one cannot expect more performance gain with the codes.

### C. Bit Shift Algorithm and BSWAP Instruction

The bit shift endian conversion codes is actually identical to `BSWAP` instruction when compiled with the optimization option of `-O2`. The disassembled codes obtained with `objdump -d` command are below:

0000000000000000 <bit\_shift>:

```

0: 48 89 f8            mov     %rdi,%rax
3: 48 0f c8            bswap  %rax
6: c3                  retq

```

## D. Endian Conversion Algorithms and Memory Alignment

It is ensured that the leading memory address (alignment) of an array is always in multiplies of 16 (16-byte alignment) in amd64 architecture. However, if one would like to read a file in multi-thread, he/she has to make the copy of the file image on memory. In such case, the alignment is not always 16-byte. Thus I performed the benchmark described in §3.4 (single thread case) but made alignment of the arraya random number.

For the benchmark, I modified `_mm_load_si128()` and `_mm_store_si128()` in the SSE2 codes into `_mm_loadu_si128()` and `_mm_storeu_si128()`, respectively, to make the codes operable. The results are summarized in Table 12.

The trend found in §3.4.1 is roughly true in this case though the all algorithms are slightly slower (within a few %) than 16-byte alignment case. Hence one does not have to get nervous about memory alignment.

## E. The Patches for CFITSIO

### E.1. fitsio2.h

```
*** fitsio2.h.org 2013-03-08 14:19:49.560538980 +0900
--- fitsio2.h 2013-01-15 14:43:10.000000000 +0900
*****
*** 96,102 ****

    #elif defined(__ia64__) || defined(__x86_64__)
        /* Intel itanium 64-bit PC, or AMD opteron 64-bit PC */
! #define BYTESWAPPED TRUE
    #define LONGSIZE 64

    #elif defined(_SX)
        /* Nec SuperUx */
--- 96,103 ----

    #elif defined(__ia64__) || defined(__x86_64__)
        /* Intel itanium 64-bit PC, or AMD opteron 64-bit PC */
! /* #define BYTESWAPPED TRUE */
! #define BYTESWAPPED FALSE
    #define LONGSIZE 64

    #elif defined(_SX)
        /* Nec SuperUx */
*****
*** 169,175 ****
```

TABLE 12  
THE ENDIAN CONVERSION TIME IN CASE THAT MEMORY ALIGNMENT IS RANDOM

Machine	Bit Shift (msec)	BSWAP (msec)	SSE2 (msec)	SSSE3 (msec)	Byte Shuffle (msec)
Machine A	413 ± 2	416 ± 3	416 ± 5	377 ± 2	3220 ± 10
Machine B	624 ± 7	641 ± 9	602 ± 9	590 ± 3	8320 ± 90

NOTE.—The endian conversion time of 423,414,686 (=29,566×14,321) double-type elements with various algorithms in case that memory alignment is random.

```

/* generic 32-bit IBM PC */
#define MACHINE IBMP
! #define BYTESWAPPED TRUE

#elif defined(__arm__)

--- 170,178 ----

/* generic 32-bit IBM PC */
#define MACHINE IBMP
! /* #define BYTESWAPPED TRUE */
! #define BYTESWAPPED FALSE
!

#elif defined(__arm__)

```

## E.2. cfileio.c

```

*** cfileio.c.org 2013-03-08 14:20:09.052539296 +0900
--- cfileio.c 2013-01-16 19:57:46.000000000 +0900
*****
*** 3763,3769 ***
    }

    /* test for correct byteswapping. */
!
    u.ival = 1;
    if ((BYTESWAPPED && u.cval[0] != 1) ||
        (BYTESWAPPED == FALSE && u.cval[1] != 1) )
--- 3763,3769 ----
    }

    /* test for correct byteswapping. */
! /*
    u.ival = 1;
    if ((BYTESWAPPED && u.cval[0] != 1) ||
        (BYTESWAPPED == FALSE && u.cval[1] != 1) )
*****
*** 3776,3782 ***
    FFUNLOCK;
    return(1);
}
!

/* test that LONGLONG is an 8 byte integer */

--- 3776,3782 ----
    FFUNLOCK;
    return(1);

```

```
    }  
! */  
  
/* test that LONGLONG is an 8 byte integer */
```



## REFERENCES

- Amdahl, G. M., Proceedings of the April 18-20, 1967, Spring Joint Computer Conference, AFIPS '67 (Spring), 483
- Eguchi, S., Kawasaki, W., Shirasaki, Y., et al. 2012, arXiv:1211.3790
- Greisen, E. W., & Harten, R. H. 1981, A&AS, 44, 371
- Grosbol, P., Harten, R. H., Greisen, E. W., & Wells, D. C. 1988, A&AS, 73, 359
- Intel Corporation, 2012, Intel 64 and IA-32 Architectures Software Developer's Manuals, Volume 2A, 3-78
- Lucas, R., Richer, J., Shepherd, D., Testi, L., Wright, M., & Wilson, C. 2004, ALMA Memo #501
- Matsumoto, M., & Nishimura, T., 1998, ACM Trans. Model. Comput. Simul., 8, 3
- Pence, W. D. 2010, Astrophysics Source Code Library, 10001
- Rohland, C., 2001, <http://www.kernel.org/doc/Documentation/filesystems/tmpfs.txt>
- Wells, D. C., Greisen, E. W., & Harten, R. H. 1981, A&AS, 44, 363